Solar Flare Prediction using Convolutional Neural Nets (37)

Kendall Johnson
kjohns21@masonlive.gmu.edu | Youtube: MoveOverRover
George Mason University | 4400 University Drive, Fairfax, VA 22030
Goddard Space Flight Lab | 8800 Greenbelt Rd, Greenbelt, MD 20771

Abstract [1]

During my time with the Space Weather Lab at George Mason University (GMU), most of our research was focused on Active Regions (AR) on the Sun's surface. Recent work with Goddard's Heliophysics Lab has opened my field to the uses of Artificial Intelligence (AI) and Artificial Neural Networks (ANN). ARs are one of the last natural phenomena that we don't fully understand what governs its movements and actions. This problem was a great fit to use an ANN algorithm to determine and decipher the qualities of the images that indicate activity when formulas and simulations fail. Knowledge of the Sun's surface and ARs are critical because, at any moment, a harmful Coronal Mass Ejection (CME) can be released causing worldwide failure of the electric grid. Fortunately, most events correlate, so when a strong solar flare occurs in an active region, it is an excellent indicator that a CME will have a stronger possibility to release from that same region. Dr. Jie Zhang, a solar physic professor and advisor at GMU, and I have recently looked at the old question of can we predict solar flares from magnetogram images of the ARs using AI? We decided that using an ANN was the most efficient approach in the fact we would be dealing with larger datasets. We attempted to train the ANN with the AR images so that when the trained ANN is presented with unknown AR images, it could correctly predict if that region will have a solar flare within 24 hours. In a combined effort with GMU's computer science department, we have now matured our ANN to a Convolution Neural Network (CNN) that is optimized for image classification. CNN is still an ANN, but it has the added feature of convolution layers that mathematical takes into account the surrounding pixels as a feature of the ANN. Convolutional layers are an excellent technique used to find structures in images using only pixel data. Our research data is the magnetogram images from Helioseismic and Magnetic Imager (HMI) on the Solar Dynamic Observatory (SDO) sliced to a square region containing the full AR. Our data is from 2010 to 2014, which consists of around 1000 images. The images are from the last solar maximum to get a more significant distribution of ARs that erupted with a solar flare within 24 hrs, and this was done by connecting them with archived flare. We are now looking toward using object detection algorithms like YOLO (you only look once) to take the entire magnetogram image of Sun to detect ARs and automatically slice them to a shape the CNN can read and predict. Our end goal is the addition of these two powerful AI techniques to produce a program that can be used by scientists and satellites to predict the release of a CME on behalf of humanity. I hope to present a proof of concept that can be used to observe the Sun's surface, and when an AR forms, the object detector will find it, and the CNN determines if a solar flare will occur within 24 hrs.

The Introduction [2]

In 1925 Cecilia Payne-Gaposchkin proposed her doctoral thesis that the stars are composed of mostly hydrogen and helium. Ever since then we have continued to learn more about our host star from the fact that is our source of life-giving energy and also the biggest, most dangerous nuclear bomb for the next 4 light-years.

The Sun continuously produces streams of charged particles into the surrounding space, which is a reason why space is dangerous. When the Sun has a lot of activity in its active regions it is likely to release a solar flare and a Coronal Mass Ejection (CME). The CME is a concentrated pressure wave of charged particles from the sun due to actions of the surface and has been known to cause problems on Earth and in our solar system. A large CME hit Earth on March 9th in 1989 causing a large geomagnetic storm shutting down some cities' power grids, like Quebec, and completely jamming worldwide communication channels like radio. It was reported that an X15 (very big) solar flare was reported on March 6th just 3 days before the incident. But, this is simply a bad day compared to the ferocity of the Carrington Event on 1859 from September 1st to the 2nd. The Carrington Event is the largest geomagnetic storm on record. Electrical grids were very small then, but the telegraph systems all over there world failed, and many telegraph workers reported that they were electrocuted at work by the event. Richard Carrington reported that a 'white light flare' came from the Sun several hours before the event. It is understood that a geomagnetic storm like the 1859 Carrington Event today would destroy electrical grids, cause widespread blackouts, and cost trillions of dollars. In 2012 a powerful CME, similar to the Carrington Event, was released but missed the Earth by nine days. In this introduction, I wanted to simply state the background and importance of this work which has led me to attempt to predict solar flare occurrences on the Sun. An accurate prediction can give us more time to be better prepared to handle it when it does.

Definitions

Active regions - Regions on the Sun's surface that have very strong magnetic fields. They have a tendency the form sunspots, seen as the darker region in Figure 1. Active regions sometimes come in contact with another polar opposite active region and will produce solar phenomena like flares and CMEs.

Charged Particles - are atomic particles or ions with an electric charge. They are released by the trillions by the Sun producing the solar wind. They can disrupt or destroy unprotected electrical equipment.

Coronal Mass Ejection (CME) - is a significant release of plasma, charged particles, and magnetic flux from the Sun's surface seen in Figure 2. They are often seen to follow the appearance of solar flares, but a CME will not always that be released with every solar flare.

Geomagnetic Storms - are a major disturbance of Earth's magnetosphere that occurs when there is a very strong exchange of energy from the solar wind to Earth's atmosphere. Can be seen on Earth as Auroras in Figure 3

Solar Flare - is a sudden bright flash of light on the Sun's surface, seen as the bright spot in Figure 2. It is found primarily seen in active regions. A solar flare can be accompanied by a CME.

Figure 1 (Magnetogram Image of AR and Solar disc)

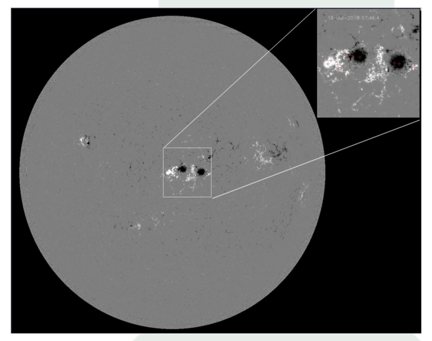

Figure 2 (Release of Solar Flare and CME)

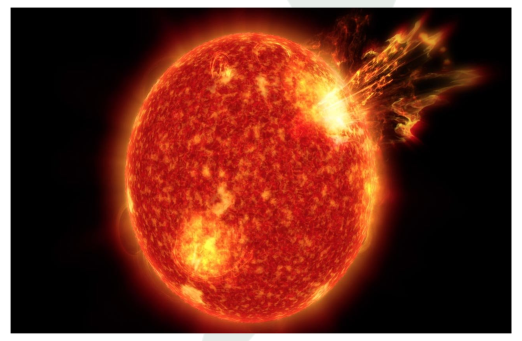

Figure 3 (Aurora from Geomagnetic storms)

Methods and Results [3]

The method I am using to predict a solar flare occurrence from an active region is the AI algorithm known as Artificial Neural Networks (ANN) and in particularly Convolutional Neural Networks (CNN). A CNN has all of the same functions and structures as ANNs except the addition of a convolutional layer that pools a set of incoming data to come up with results that take into account the values of close proximity inputs. This can be done by simply finding the sum or the average of pieces of the incoming inputs. CNNs are known to work very well with image data that is in the matrix-like form similar to pictures from a camera.

The CNN architecture I am using started with the VGG-16 architecture, seen in Figure 5, which is a very successful CNN algorithm used by the Visual Geometric Group out of Oxford who got a 95.2% model score on the very large ImageNet set of images. VGG-16 had a normal 224 x 224 input but I needed 256 x 256, so I made an hourglass-like residual layer structure to start the network, an example is seen in Figure 4 added to the front of Figure 5. I did this to strengthen the starting heat map of features and output to the smaller 224 x 224 image size when added to the beginning of the VGG 16 model. I also removed Max-pooling and used average pooling due to the way the numeric data was presented. I used the activation function ReLu for every layer but the last which I used the softmax function to help classify. My loss function was a simple Binary crossentropy with no reduction type and my optimizer was Adam with a learning rate of 0.0001, amsgrad = True, beta 1 = 0.9, beta 2 = 0.9999, and epsilon = 0.00001.

Figure 4 (Hourglass ANN architecture)

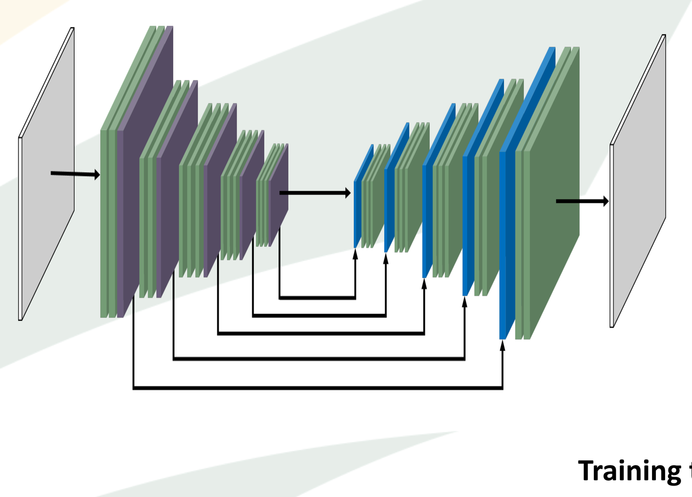

Figure 5(VGG-16 ANN architecture)

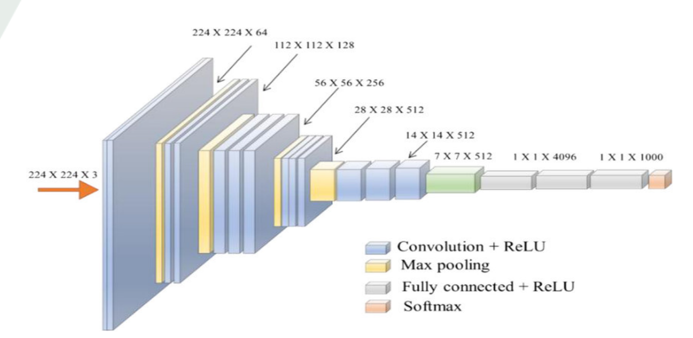

Training the Model

The data used to train the model was 1070 images of the Sun's active regions, like in Figure 1, with the dimensions of 256x256x1. The images are HMI Magnetogram, like in Figure 1, were taken from the Solar Dynamic Observatory (SDO) from 2010 to 2014 around the last solar maximum. There were 70 images of active regions that created a solar flare within 24hrs, and 1000 images of active regions that did not. I also used 22,000 epochs with an image batch size of 32. My goal was simply to produce a 0 meaning that no flare will be produced or a 1 meaning that a solar flare will be released.

The Hardware that is used to train the following CNN is two NVIDIA Quadro RTX 4000 GPU for a total of almost 5,000 Cuda cores. The software programs used the create and train the CNN architecture are TensorFlow 2 and Keras in python 3.7 on the Ubuntu 18.04 operating system.

From the trained CNN model I produce the visual results of a confusion matrix (Figure 6), a normalized confusion matrix (Figure 7), and a ROC curve (Figure 9), and for the quantitative model parameters produced are Sensitivity, Specificity, Precision, Negative Predictive Value, Accuracy, AUC (area under the curve), and the Appleman Skill Score (Equation 1). The calculations are shown in Figure 8 and the calculated values are shown in Table 1. I converted some to a percentage for discussion purposes.

Figure 6 (Confusion Matrix)

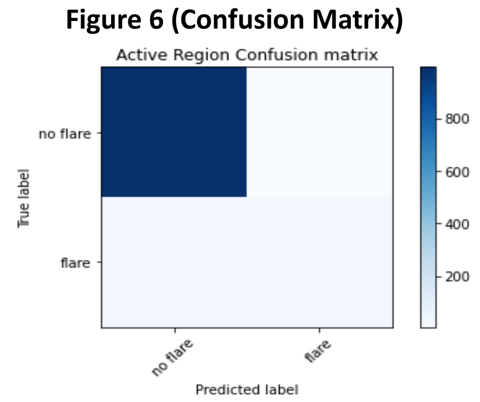

Figure 7 (Normalized Confusion Matrix)

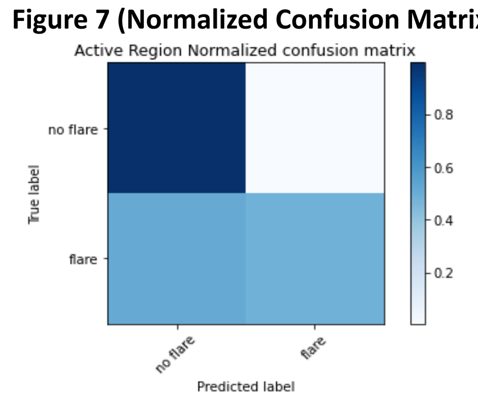

Figure 9(ROC Curve)

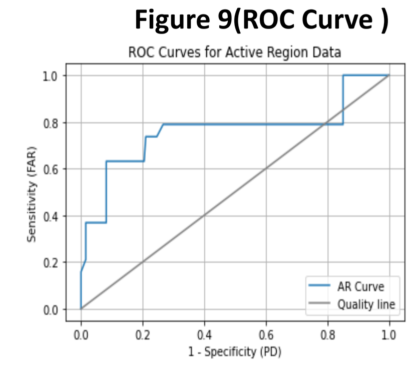

Figure 8 (Confusion Matrix labels and Equations)

		Predicted Class: Positive	Predicted Class: Negative	
Actual Class	Positive	True Positive (TP)	False Negative (FN) Type II Error	Sensitivity $\frac{TP}{(TP+FN)}$
	Negative	False Positive (FP) Type I Error	True Negative (TN)	Specificity $\frac{TN}{(TN+FP)}$
		Precision $\frac{TP}{(TP+FP)}$	Negative Predictive Value $\frac{TN}{(TN+FN)}$	Accuracy $\frac{TP+TN}{(TP+TN+FP+FN)}$

Table 1 (Calculated Vales from CNN Model)

Model Score	92.21%
Sensitivity	48.57%
Specificity	99.40%
Precision	85%
Negative_Preditive	96.50%
Accuracy	96.07%
Applemen Skill Score	-0.03%
AUC	.761

Equation 1 : $$Appleman's\ Skill\ Score = \frac{\frac{(TN+TP)}{total\ images} - (TN+FP)/total\ images}{1-(TN+FP)/total\ images}$$

Methods and Results (continued..)

Picking Active regions from the Sun's disc.

I attempted to make an object detection model that can detect and classify active regions on the sun, but most, if not all, of the activity of the Sun, will have an active region present. What I did instead is use the python module OpenCV to create a sudo-edge detector. I used Canny edge detection, gaussian blurs, and threshold adjustments for the program to detect ARs on the surface against the seeming blank surface background the magnetogram produces. When an AR is detected a bounding box is put around it, seen in Figure 10, and this bounding box is what will in the future slice to a specific size image and feed into the train CNN. I believe this only worked because of simplistic magnetogram images and would work poorly for any other solar image of wavelength.

Figure 10 (AR detections on the Solar Disc)

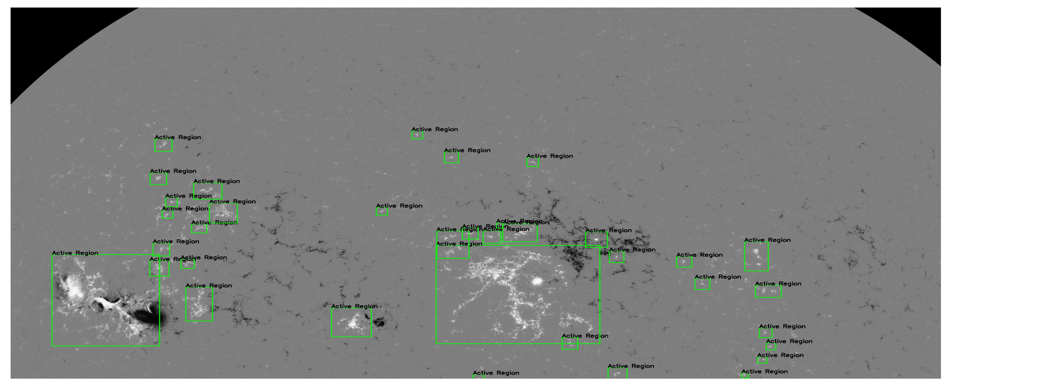

Discussion and Conclusion [4]

Discussion

The model performed well with a respectable model score of 92.2%, but this does not mean that the CNN can predict the occurrence of solar flares at that efficiency. First, in discussing the convolution matrix in Figure 6 it is easy to see that the data is completely skewed by the overwhelming higher percentage of non-solar flare making active regions by the fully blue true positive and every other box being almost blank. For the normalized confusion matrix seen in Figure 7, we can see the distribution of squares better, but we are also able to see the poor classifying for active regions that will produce a solar flare from the false positive and true negative beginning about the same color. This says that it was around a 50/50 split of the trained CNN model being able to determine if the region that is going to have a solar flare will have a solar flare. Just from the confusion matrixes, we can say that the CNN is good at determining that an active region won't have a solar flare, but we are likely to get quite a lot of false positives before getting a true solar flare.

The ROC curve, in Figure 9, agrees with the result of getting very many false positives with the slope line closer to a lower specificity on the x-axis. The ROC curve's AUC shows positive predictability with .761 that represents moderate model performance.

All calculated predictive values, except specificity and Appleman's skill score, are very good but are taken with the same grain of sand the great model score is. This is because the model was very good a predicting the non-flare bearing active regions that make up more than 90% of the data. The specificity of 45% on the other hand shows the truth of not being able to identify the flare bearing active regions results very well, and this reiterates the runaway false positive problem above.

The Appleman's skill score was my most sought after metric to truly quantify how well my model worked. In the 2016 paper titled A COMPARISON OF FLARE FORECASTING METHODS, G.Bares et al. were able able to get a state-of-the-art model with an Appleman's skill score of 0.19. simply meaning they were able to predict more solar flares than not. My Appleman's skill score was -0.03, basically zero, because of the 50/50 tie between false positives and true negatives of flare bearing active regions. This result doubles down on the fact that we will get detect a similar amount of false positives of flare-bearing active regions as actual solar flare events.

Conclusion

In conclusion, it is no simple feat to predict solar flares. Although there were many positive results in this proof of concept the algorithm is far from being a great predictor of the occurrences of solar flares. The two biggest problems are that there is such a big difference in the quantity of data for each class, and the classic problem of we need more data. The positive progress of both of these problems of both these problems, it will create a better predictive model. But, the model we created model still has the strong ability to tell that an active region is not going to have a solar flare. This model may not have certainty in that a possible active region may have a solar flare, but due to the vast amounts of false positives, it will likely not miss the actual AR that will birth a solar flare. Much like the algorithm to find fraudulent credit card activity, it is better to find the problem and be wrong about that finding than to miss the problem entirely. One of the many false positives this trained CNN would detect could be the real thing, and that may be the difference in readiness for a Carrington-like solar event.

References [5]

(1) G.Barnes et al. (2016) A COMPARISON OF FLARE FORECASTING METHODS. I. RESULTS FROM THE"ALL-CLEAR"WORKSHOP
(2) Karen Simonyan+& Andrew Zisserman, (2015), VERY DEEP CONVOLUTIONAL NETWORKS FOR LARGE-SCALE IMAGE RECOGNITION
(3) Nushaine Ferdinand (May 29, 2020), Using Hourglass Networks To Understand Human Poses

Acknowledgements [6]

I would like the thank all the great staff of GMU's Physic's department, College of Science Research Club, Robotic Lab, The Space Weather Lab, along with the great staff at NASA Goddard's Heliophysic lab. I would like to particularly thank Jie Zhang for his idea to do this work, Barbara Thompson for her help and support with AI, and Suman Dhakal for the hard task of getting and labelling the data.